Accepted by ApJ Oct 23, 2016
Preprint typeset using LaTeX style emulateapj v. 5/2/11

arXiv:1606.09232v2 [astro-ph.HE] 25 Oct 2016

NEUTRON STAR MASS-RADIUS CONSTRAINTS USING EVOLUTIONARY OPTIMIZATION

A. L. Stevens[1,2], J. D. Fiege[3], D. A. Leahy[4], and S. M. Morsink[1]

Accepted by ApJ Oct 23, 2016

ABSTRACT

The equation of state of cold supra-nuclear-density matter, such as in neutron stars, is an open question in astrophysics. A promising method for constraining the neutron star equation of state is modelling pulse profiles of thermonuclear X-ray burst oscillations from hotspots on accreting neutron stars. The pulse profiles, constructed using spherical and oblate neutron star models, are comparable to what would be observed by a next-generation X-ray timing instrument like ASTROSAT, NICER, or LOFT. In this paper, we showcase the use of an evolutionary optimization algorithm to fit pulse profiles to determine the best-fit masses and radii. By fitting synthetic data, we assess how well the optimization algorithm can recover the input parameters. Multiple Poisson realizations of the synthetic pulse profiles, constructed with 1.6 million counts and no background, were fitted with the Ferret algorithm to analyze both statistical and degeneracy-related uncertainty, and to explore how the goodness-of-fit depends on the input parameters. For the regions of parameter space sampled by our tests, the best-determined parameter is the projected velocity of the spot along the observer's line-of-sight, with an accuracy of $\leq 3\%$ compared to the true value and with $\leq 5\%$ statistical uncertainty. The next best-determined are the mass and radius; for a neutron star with a spin frequency of 600 Hz, the best-fit mass and radius are accurate to $\leq 5\%$, with respective uncertainties of $\leq 7\%$ and $\leq 10\%$. The accuracy and precision depend on the observer inclination and spot co-latitude, with values of $\sim 1\%$ achievable in mass and radius if both the inclination and co-latitude are $\gtrsim 60°$.

Subject headings: stars: neutron — stars: rotation — X-ray: binaries — relativity — pulsars: general — methods: numerical

1. INTRODUCTION

Neutron stars are an astrophysical laboratory for studying cold supra-nuclear-density matter. Accreting millisecond X-ray pulsars, a particular subset of neutron stars in low-mass X-ray binaries (LMXBs), are rapidly spinning accretion-powered neutron stars with spin periods of a few milliseconds (e.g., SAX J1808.4−3658, Wijnands & van der Klis 1998). Their pulsed X-ray emission originates from material striking the surface of the neutron star during regular accretion and warming an area on the surface so that it emits blackbody radiation. Then, as the neutron star rotates, it gives periodic oscillations in brightness as the emitting region faces towards and away from the observer. Since these photons originate from the surface of the neutron star itself, physical properties like its mass and radius are encoded in the detected pulse profile. Fitting these pulse profiles with realistic models can then yield neutron star mass and radius estimates (Watts et al. 2016).

In addition to regular pulsed X-ray emission, some neutron stars in LMXBs exhibit thermonuclear (Type I) X-ray bursts (Watts 2012) . In a fraction of thermonuclear X-ray bursts, we observe brightness oscillations, where the frequency corresponds strongly with the spin frequency of the neutron star; these are referred to as thermonuclear burst oscillations. The pulse profile models that we discuss in this paper refer specifically to models of these burst oscillations.

The emission area on the surface of the neutron star is referred to as the hotspot or spot. Theories suggest two different surface hotspot models: one that ignites nuclear burning at one point and spreads across the whole neutron star, and another that ignites at one point and begins to spread but remains limited to a smaller area (referred to as a "persistent hotspot") (see Watts 2012 and references therein). The persistent hotspot on the surface of a rotating neutron star has been demonstrated to be an effective model for Type I X-ray bursts from the source 4U 1636−536 (Artigue et al. 2013), and so we use a persistent spot model with no size variation over the course of the burst. The fixed spot model is used for convenience here; a changing spot model could be incorporated for observations that show evidence for such behaviour.

The spectral model depends on the physics of the spot production and includes both the energy and angular dependence of the emitted radiation. In the case of rotation-powered X-ray pulsars (Bogdanov 2013), a hydrogen atmosphere model is appropriate (e.g., Heinke et al. 2006). The hydrogen atmosphere model depends on the spot's temperature and the local surface gravity. Since the surface gravity only depends on the mass, radius, and spin of the star (AlGendy & Morsink 2014), the local temperature is the only additional free parameter introduced by the spectral model.

For accretion-powered X-ray pulsars, an empirical model including a blackbody plus Comptonized photons has been used, motivated by spectral observations (Poutanen & Gierliński 2003; Leahy et al. 2009, 2011; Morsink & Leahy 2011).

[1] Department of Physics, University of Alberta
[2] Anton Pannekoek Institute for Astronomy, University of Amsterdam
[3] Department of Physics and Astronomy, University of Manitoba
[4] Department of Physics, University of Calgary

The Comptonization model includes photon power law indices as well as a parametrized fan-beaming model. These models required two free parameters, one for the energy dependence and the second for the angular dependence of the emitted radiation. One of the issues seen with fitting the data from the accretion-powered X-ray pulsars is that the extra degree of freedom in the radiation's angular dependence leads to extra degeneracies amongst the geometric parameters. The result is fairly large regions of parameter space allowed by the fits, which do not strongly constrain the neutron star's equation of state.

Since thermonuclear X-ray burst oscillations can be well-reproduced with a simpler spectral model, we can fit models in a reduced but still physically motivated parameter space. Fewer free parameters gives fewer degeneracies among the parameters, and therefore allows for better constraints on the neutron star's mass and radius.

In the setup of our models, we assume that the inner boundary of the accretion disk is the neutron star's co-rotation radius (Ghosh & Lamb 1979). From this assumption, the accretion disk would block emission from a possible second antipodal hotspot (from an observer's perspective), so we only test models for one spot in the northern hemisphere of the neutron star. If the signature of an antipodal spot is detected, our code can be easily adapted to include a second hotspot.

We created synthetic data for a variety of different neutron star and spot parameters. Fitting multiple Poisson realizations for each synthetic pulse profile allows us to disentangle what uncertainty is due to random statistical fluctuations and what is due to inherent degeneracy between the parameters; understanding both is crucial for placing proper constraints on neutron star masses and radii. The pulse profile fitting was carried out with the Ferret optimization algorithm (Fiege 2010) to determine the acceptable range of masses and radii.

In this paper we show that evolutionary optimization algorithms are a powerful method of fitting neutron star pulse profiles, and we test the effects of changing various input parameters on how well we can recover the true neutron star mass and radius. In Section 2 we explain the details of constructing the pulse profiles, our parameter choices, and the Ferret algorithm. The results of the pulse profile fitting are described and examined in Section 3, and the conclusions are discussed in Section 4.

2. PULSE PROFILE MODELS

We construct the pulse profiles within the Schwarzschild + Doppler (S+D) approximation (Miller & Lamb 1998; Poutanen & Gierliński 2003) and the Oblate Schwarzschild (OS) approximation (Morsink et al. 2007). In the S+D and OS approximations, the metric exterior to the rotating neutron star is approximated by the Schwarzschild metric as described by Pechenick et al. (1983), adding the appropriate Doppler boost factors arising from the rotation of the star. In the S+D approximation, the surface of the star is a sphere, while in the OS approximation the surface is an oblate spheroid. Cadeau et al. (2007) compared the S+D and OS results with pulse profiles generated from the exact metric and showed that the OS approximation is a good approximation for stars spinning with frequencies above 300 Hz. However, we continue to use the S+D approximation in many of our models in order to further explore the effect of using the wrong shape on the fits. At spin frequencies higher than 600 Hz, it may be necessary to use higher order approximations that make use of the star's quadrupole moment (Psaltis & Özel 2014), however this level of approximation is not necessary for the stars studied in this paper.

Pulse profiles can be constructed once 8 geometric parameters and a spectral emissitivity model are specified. The 8 geometric parameters are the neutron star's mass M; equatorial radius R; spin frequency $\nu_{\rm spin}$; the observer's inclination angle i (as measured from the spin axis); the hotspot's co-latitude θ; the angular radius of the spot ρ; the distance to the star d; and a phase offset ϕ. In practice, the star's spin frequency will always be known, so there are only 7 geometric parameters. It is possible to add parameters describing a more complicated shape for the spot (Poutanen et al. 2009), but in this paper we only consider the simplest spot models, which are circular spots with uniform temperature. Due to the approximately universal nature of a spinning neutron star's shape (Morsink et al. 2007; Bauböck et al. 2013), inclusion of the star's oblate shape does not require any additional free parameters.

Thermonuclear X-ray burst oscillations can be spectrally modelled as a single-temperature blackbody with limb-darkening (Artigue et al. 2013). Once a prescription for the angular dependence has been selected, such as the Chandrasekhar (1960) limb-darkening model (approximated by the Hopf function), the only free parameter is the hotspot's temperature. The reduced parameter space required by this spectral model results in less degeneracy with the geometric parameters. For this reason, the oscillations from Type I X-ray bursts will be a major target for large-area, high-time-resolution X-ray telescopes like ASTROSAT (Singh et al. 2014), the soon-to-be launched NICER mission (Arzoumanian et al. 2014), and a future LOFT-like mission (Feroci et al. 2012).

The geometric parameters M, R, i, and θ have inherent degeneracies, so it is useful to refer to less-degenerate combinations of these parameters. In the spherical S+D approximation, i and θ only appear in the formulae in the combinations $\sin i \sin\theta$ and $\cos i \cos\theta$. As a result, in all fits there is a simple degeneracy that allows i and θ to be switched. Likewise, M and R are somewhat degenerate, so the dimensionless compactness ratio M/R can be better-constrained. These parameter combinations factor into the approximate bolometric pulse amplitude Amp (Beloborodov 2002),

$$\mathrm{Amp} = \frac{(1-2M/R)\sin i \sin\theta}{2M/R + (1-2M/R)\cos i \cos\theta}, \tag{1}$$

and the dimensionless projected velocity of the spot β,

$$\beta = \frac{2\pi R \nu_{\rm spin}}{\sqrt{1-2M/R}} \sin i \sin\theta, \tag{2}$$

TABLE 1
Summary of Test Models

Name	Description	Shape	$\nu_{\rm spin}$ Hz	M $M_\odot$	R km	i deg.	θ deg.	ϕ	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β
A	Fiducial	Sphere	600	1.6	12	60	20	0	0.197	0.296	0.470	0.373	0.057
AO	Oblate	Oblate	600	1.6	12	60	20	0	0.197	0.296	0.470	0.373	0.057
θ_{37}	$\theta = 37°$	Sphere	600	1.6	12	60	37	0	0.197	0.521	0.399	0.720	0.101
θ_{37}O	$\theta = 37°$	Oblate	600	1.6	12	60	37	0	0.197	0.521	0.399	0.720	0.101
θ_{60}	$\theta = 60°$	Sphere	600	1.6	12	60	60	0	0.197	0.750	0.250	1.305	0.145
$(M/R)_{hi}$	High M/R	Sphere	600	1.68	11	60	20.8	0	0.226	0.308	0.467	0.350	0.057
$(M/R)_{lo}$	Low M/R	Sphere	600	1.35	13	60	19.8	0	0.153	0.293	0.471	0.423	0.057
$\beta_{hi}(M/R)_{hi}$	High β, high M/R	Sphere	600	1.68	11	60	64	0	0.226	0.778	0.219	1.235	0.145
ν_{400}	Low $\nu_{\rm spin}$	Sphere	400	1.60	12	60	20	0	0.197	0.296	0.470	0.373	0.038

in geometric units ($G = c = 1$). Due to the reduced degeneracy, it is possible to fit for $\sin i \sin\theta$, $\cos i \cos\theta$, M/R, Amp, and β better than the individual parameters. In our models, we consider an infinitesimally small spot for simplicity. However, the $i - \theta$ degeneracy can be partially broken for models with a large spot, in which case the spot would span a range in θ but not in i.

Equation (1) is only an approximate relation for the bolometric pulse amplitude. In reality, the pulse amplitude depends both on the emitted energy spectrum and the energy bands at which the observations are made. There is no simple formula for the dependence of the pulse amplitude on energy, but it can be computed numerically. Since the pulse amplitude Amp depends on the observed photon energy, observations of the burst oscillations in two or more energy bands can provide stronger constraints than suggested by equation (1); for this reason, two energy bands are used in this work. Additional information could be extracted with more energy bands, but we were unable to accommodate more bands. The projected velocity β controls the asymmetry in the pulse profile through the Doppler boosting effect, as well as the phase lags between the hard and soft energy bands. For higher values of β, the pulse profiles are more asymmetric in rise and fall times, showing the effects of higher harmonics.

Previous work (Morsink & Leahy 2011; Stevens 2013; Lo et al. 2013; Bauböck et al. 2015) has shown that for smaller spots, it is sufficient to compute the pulse profile assuming a point-source spot instead of an extended region. By adopting the point-source approximation, we simplify the calculation and do not make use of the angular radius parameter ρ, reducing the number of free geometric parameters to 6. Furthermore, we normalize the synthetic pulse profiles to have a mean of 1 to remove the dependence on the distance d, giving five free geometric parameters for each model.

Our goal is to constrain the neutron star's mass and radius based on fitting models to the pulse profile, and determine how the shape of the pulse profile affects how well we can constrain its M and R. By fitting synthetic data, the input parameters are known, so we can analyze how well the fitting can recover the true M and R. In this paper we investigate two sources of error in the constraints: degeneracy-related uncertainty and statistical uncertainty. Although the two types of uncertainties are coupled we have introduced two measures of the uncertainties $\sigma_{\rm degen}$ and σwhich are affected differently by the degeneracy and the statistics. The degeneracy-related uncertainty $\sigma_{\rm degen}$ arises from parameter degeneracies in each fit, whereas the standard deviation σ is examined by simulating multiple Poisson realizations of a model and determining the mean and standard deviation over all the fits. We explore how σ and $\sigma_{\rm degen}$ are affected by the values of Amp and β for each synthetic pulse profile. We also quote the accuracy of the fit in M and R, comparing the mean best-fit value for each model with the true value.

2.1. *Properties of Test Models*

We computed pulse profiles for a set of nine test models with the parameter values listed in Table 1. The properties that were kept the same for all test models are the spot's temperature, the phase offset ϕ, the observer's inclination angle i and the distance to the star d. The spot emission model is a 2 keV blackbody (in the frame comoving with the neutron star's surface) with a limb-darkening atmosphere, approximated by the Hopf function (Chandrasekhar 1960), appropriate for a Thompson-scattered atmosphere. Since for real data we would independently have the spot temperature *at infinity* instead of in the star's comoving frame, this parameter should be allowed to vary within a narrow range (explored in Section 3.7), where the appropriate range would be determined by the temperature at infinity. The model assumes that any emission from the surface of the neutron star outside the spot is negligible. The observers inclination angle $i = 60°$ and the phase offset $\phi = 0$ for all cases. The computed pulse profiles are normalized to an average flux of 1, so that the distance d to the star does not affect the pulse profiles.

The parameters that were changed for different test models are M, R, θ, the star's shape (spherical or oblate), and $\nu_{\rm spin}$. For oblate neutron star models, R is defined to be the radius of the neutron star at the spot. The formalism for the oblate model is detailed in Morsink et al. (2007), and does not require the addition of any extra parameters.

The parameters for the fiducial model, Model A, were chosen to be representative of the masses and radii of accreting millisecond neutron stars. Our fiducial mass $M = 1.6\ M_\odot$ is larger than the mass $M = 1.4\ M_\odot$ typically measured for slowly rotating radio pulsars, since we expect the neutron star has been spun up by and gained mass from accretion. The radius of 12 km is consistent with other radius estimations (Leahy 2004; Steiner et al. 2010). Rapid rotation with $\nu_{\rm spin} \geq 550$ Hz is seen in burst oscillations for at least seven neutron stars that exhibit Type I X-ray bursts (Watts 2012). We have chosen $\nu_{\rm spin} = 600$ Hz for the fiducial model since it is representative of these rapid rotators. The angles

were chosen to provide a pulse amplitude (Amp = 0.373) comparable to the largest pulse amplitudes seen, in order to reduce the parameter degeneracy. In particular, the neutron star 4U 1636-536 has rms pulse amplitudes as high as 0.25 after subtracting the pre-burst emission (Strohmayer et al. 1998; Galloway et al. 2008), where Amp = $\sqrt{2}$ rms. The full set of theoretical parameters are given in the row labelled "A" in Table 1. The resulting pulse profiles in two energy bands are shown in Figure 1 with black and red solid curves labelled "True". Our fiducial model is similar to the model described as "low inclination" by Lo et al. (2013), except that they consider a slightly smaller R and a slower $\nu_{\rm spin}$. As will be shown in Section 3.5, small changes in R do not qualitatively change the results of this paper. Our choice of a faster $\nu_{\rm spin}$ improves the accuracy of the fits, but the faster $\nu_{\rm spin}$ is also more appropriate for the neutron stars that we hope to apply this method to.

Theoretical pulse profiles were calculated for the nine different test models. Each theoretical pulse profile was converted to a set of twenty synthetic observations by adding noise from a Poisson distribution to the pulse profile. The standard case (a low-noise model) assumes 25,000 photon counts per phase bin and no background count rate, as if the background were negligible in comparison with the hot spot emission. By not including a background, we are underestimating the Poisson fluctuations, and thus underestimating the error bars. This is done so we can test the suitability of evolutionary optimization on the best possible quality of synthetic data.

For a variation on the model θ_{60}, a higher noise level was used: $\theta_{60}C_{6250}$ assumes 6250 photon counts per phase bin with no background. In total, there are ten sets of twenty simulated observed pulse shapes.

Each of the 200 simulated pulse profiles was then fit using the Ferret algorithm (described in the next subsection) which searches for different values of the free parameters in order to minimize the χ^2 fit statistic. Five parameters were allowed to vary within physical ranges: $1.0\,{\rm M}_\odot \le M \le 2.5\,{\rm M}_\odot$, $6.0\,{\rm km} \le R \le 16.0\,{\rm km}$, $0^\circ \le (i,\,\theta) \le 90^\circ$, and ϕ defined cyclically from 0 to 2π. The star's shape (spherical or oblate) was fixed to be the same in the fitting procedure as in the theoretical waveform. The temperature is kept fixed at 2 keV since it has been shown by Lo et al. (2013) that observations in multiple energy bands allows for a good determination of the gravitationally-redshifted spot temperature. We do not allow the spot size to vary, as explained in the previous subsection. For applications to real data, it would not be difficult to add variations in spot size (and shape), temperature, and to use unnormalized fluxes to allow for a distance measurement. However, the addition of the extra free parameters for the number of synthetic waveforms considered in this study would be impractical. We carried out preliminary trials (summarized in Section 3.7) allowing these parameters to vary but did not find much change in our final results.

Since there are two energy bands with 32 phase bins each, and five free parameters, there are 59 degrees of freedom (dof). The fitting yielded best-fit model parameters and also allowed computation of confidence regions in the M-R plane, from which we can assess the effect of various parameters on the uncertainties in parameter determination.

2.2. *Parameter Fitting using Evolutionary Optimization*

Evolutionary optimization algorithms provide a useful but heuristic approach to search through large parameter spaces, to optimize the fit of a model to data. Such algorithms use principles inspired by biology to evolve a population of candidate parameters sets, over many generations, toward an optimal solution. The heuristic nature of the search aims to sample the parameter space efficiently, but not completely, since exhaustive search is not practical for problems involving many parameters. Therefore, it cannot be guaranteed that the true, globally optimal solution has been found, although this would also be true for any other optimization algorithm, besides exhaustive search. Evolutionary algorithms have been well studied and found to be useful in many fields of science and engineering.

The oldest, and most commonly known, type of evolutionary optimizer is the genetic algorithm (Holland 1975; Goldberg 1989, 2002). Classic genetic algorithms encode their search parameters on a (usually) binary string (genotype), which is decoded into a model (phenotype). A population of candidate parameter sets is normally initialized as a set of random bit strings, which are expressed as a model, and in turn evaluated by a fitness function. This information is used to probabilistically select good parameter sets (individuals) to propagate to the next generation, in analogy to the "survival of the fittest" principle of Darwinian evolution. Parameter sets undergo random bitwise mutations to help perturb solutions into previously unexplored regions of parameter space. Information is shared between individuals by means of a crossover operator that cuts a pair of bit strings at a random position, and recombines them into a new "offspring" configuration, in analogy with sexual reproduction. This process of evaluation, selection, mutation, and selection is performed iteratively, over many generations, until a convergence criterion terminates the search. The genetic algorithm can be viewed as a directed stochastic search, which makes use of random noise to evade local minima, while a mostly deterministic selection operator pushes solutions toward an optimal solution.

In this paper, we used versions 5.3-5.5 of Ferret from the Qubist Global Optimization Toolbox for MATLAB (Fiege 2010; Rogers & Fiege 2011), a commercially available software package, to find the best fits to our synthetic data. This is an alternative to the Markov chain Monte Carlo approach in Lo et al. (2013) to fit pulse profiles. Ferret's development began in 2002, as a variant of the multi-objective genetic algorithm, which made use of a real-valued parameter encoding and real-valued mutation and selection operators, rather than the traditional binary encoding discussed above. Numerous other features have been added to Ferret since then, which go beyond the usual genetic algorithm paradigm. Most notably, the code contains a unique linkage-learning algorithm, which detects a certain type of non-linearity (linkage) between parameters, with the goal of dividing a large parameter space into several smaller and nearly independent subspaces, thus greatly simplifying the problem. This capability is discussed at length in the software user's manual (Fiege 2010). Ferret also contains an algorithm that allows several of its most important control parameters, including the typical strength of mutation and crossover events, to be automatically adapted and optimized

during a run. This auto-adaptation capability is similar in spirit to the self-adaptation used in Evolution Strategies (ES) codes, although ES codes do not typically employ a crossover operator. Ferret does not adhere to the strict definition of a genetic algorithm, due to these enhancements and others, but the code still remains closer to the genetic algorithm paradigm than to others within the family of evolutionary optimizers.

As an example of Ferret's inner workings, consider the parameters i and θ. Ferret selects sets of values from the allowed region of parameter space ($0.01^\circ \leq i, \theta \leq 90^\circ$) for the population of a generation, and fits pulse profile models. It then keeps the best-fitting parameter sets and selects new sets to explore more of the parameter space. Within a few generations, Ferret discovers the degeneracy between i and θ, since their values can be swapped with minimal impact on the χ^2 fit. A more detailed explanation of Ferret in relation to pulse profile modelling can be found in Stevens (2013) Chapter 4. Evolutionary optimization algorithms have also been used in gravitational lensing (Rogers & Fiege 2011; Rogers & Fiege 2012), medical physics (Fiege et al. 2011), star formation (Franzmann 2014), and X-ray spectral fitting (Rogers et al. 2015) applications.

Our problem is quite easy for Ferret, which consistently finds the global minimum χ^2 fit statistic within a few generations. We note that the true global minimum is known, since our results are based on artificial data tests. After finding the χ^2 minimum, the algorithm was allowed to run for approximately 100 generations, as Ferret accumulated solutions within the confidence region. No specific convergence criterion was implemented; rather, we terminated the run manually when it was evident from the software's graphical user interface that no further improvements were being made to the minimum χ^2, and enough solutions had been accumulated within the confidence region to make contours maps of acceptable quality. Ferret was executed in MATLAB R2011a on 12-core AMD Linux servers (dual socket Opteron 2439 SE, with 32Gb RAM, running Red Hat version 4.1.2), using Ferret's built-in parallel computing features. Each fit took approximately 8-20 hours.

3. RESULTS

Pulse profiles were simulated as a set of Poisson realizations of a theoretical pulse profile model with known input parameters (see Section 2.1). Each Poisson realization was fit with Ferret to produce a set of best-fit parameters for the pulse profile along with confidence regions for the parameters. Since each theoretical pulse profile has multiple Poisson realizations, the average and standard deviation for the best-fit parameters of each theoretical pulse profile are computed.

Here we discuss the fit results of the pulse profiles, which illustrate the effects of changing different system properties. Table 2 shows a summary of the fits to the 20 different Poisson realizations for each of the input models. The first row of each pair of rows shows the input parameters of the theoretical pulse shape model. The second row shows the means and standard deviations of the best-fit values from the fits to the 20 different Poisson realizations of each model. In addition to the fit parameters (M, R, i, θ, and ϕ), other useful measures of the model (M/R, $\sin i \sin\theta$, $\cos i \cos\theta$, Amp, and β) are given.

A brief overview of the entire set of fits is given first. Some parameters are well-determined and others poorly-determined. The standard deviation of any given parameter for each set of Poisson realizations is given as the error σ of that parameter. We take the difference between the mean of a parameter for each set and the input parameter as a measure of the accuracy of that parameter fit. The means and standard deviations are listed in Table 2, from which the percent error for precision and accuracy can be determined. The Ferret algorithm also computes contour regions as a measure of the degeneracy-related uncertainty $\sigma_{\rm degen}$. The σ errors are sometimes smaller than the 1-$\sigma_{\rm degen}$ region as shown in figures in the following subsections. These two values of error are measures of the degeneracy ($\sigma_{\rm degen}$) and the statistical fluctuations (σ), but they need not be the same. The 1-$\sigma_{\rm degen}$ contours determined by Ferret are models with different i, θ, and ϕ parameters that all provide an equivalently good fit for that M and R. In this way, the 1-$\sigma_{\rm degen}$ limit can be thought of as a measure of the degeneracy of the parameter space near the best-fit model. The standard deviations σ computed from the ensemble of Poisson realizations provide a measure of the error from the statistical fluctuations. However, the strong parameter degeneracy is certainly still an important factor in the standard deviation computation, and as such the reported σ values cannot be assumed to be purely statistical error.

We find that the projected velocity β is generally the most accurate and precise, with accuracy of $0.5-3\%$ and precision of $1-6\%$. M and R are the next best, with with accuracy of $1-8\%$, but R has worse precision ($1-21\%$) than M ($1-13\%$). Compactness (M/R), $\sin i \sin\theta$, and Amp all have similar accuracies ($\sim 1-8\%$) and precisions ($\sim 1-16\%$). The accuracy of $\cos i \cos\theta$ is $\sim 4-30\%$ and precision is $\sim 2-20\%$. The least-well-determined are i and θ, with accuracies of $\sim 4-30\%$ and precisions of $\sim 5-50\%$ due to the degeneracy in angles.

In the following subsections we discuss in detail the results from fitting the fiducial model A, and then discuss trends we find by comparing the fits from model A with the fits for the models with modified parameters.

3.1. *Model A: Fiducial Model*

Twenty realizations of the Poisson noise were made to simulate observations of the fiducial Model A. The resulting flux values with error bars for the first Poisson realization of the data are shown as black and red points labelled "A1" on Figure 1. The data points with Poisson errors were used as input to Ferret, which searched for the best fit to the A1 dataset by minimizing χ^2. In the case of the A1 dataset, the pulse profiles that best fit the data are displayed with blue dashed and dotted curves in Figure 1. The parameters corresponding to the best fit of the A1 dataset ($\chi^2 = 58.2$) are shown in the row labelled "1" in Table 3. For reference we have included the derived quantities M/R, $\sin i \sin\theta$, $\cos i \cos\theta$, Amp, and β to this table.

TABLE 2
Summary of Statistical Properties of Fits

Model χ^2(59 dof)	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times10^{-2}$
A 59.3±7.6	1.60 1.52±0.11	12.0 11.4±1.2	60.0 34.7±15.0	20.0 41.8±15.2	0.0 2.8±4.6	0.1969 0.199±0.016	0.296 0.318±0.032	0.470 0.549±0.120	0.373 0.363±0.025	5.74 5.81±0.20
AO 59.1±7.0	1.60 1.65±0.06	12.0 11.7±0.9	60.0 49.3±18.2	20.0 28.3±11.9	0.0 0.8±5.3	0.1969 0.209±0.016	0.296 0.299±0.019	0.470 0.511±0.142	0.373 0.353±0.060	5.74 5.76±0.18
θ_{37} 61.1±9.9	1.60 1.62±0.08	12.0 12.4±0.9	60.0 49.1±13.8	37.0 47.9±13.0	0.0 -0.7±2.9	0.1969 0.194±0.009	0.521 0.509±0.033	0.399 0.386±0.064	0.720 0.726±0.029	10.10 10.09±0.17
θ_{37}O 58.9±8.1	1.60 1.58±0.04	12.0 11.8±0.4	60.0 49.8±9.7	37.0 46.3±7.8	0.0 1.1±1.5	0.1969 0.198±0.008	0.521 0.528±0.016	0.399 0.424±0.029	0.720 0.704±0.037	10.10 10.08±0.14
θ_{60} 57.7±8.4	1.60 1.61±0.01	12.0 12.1±0.1	60.0 61.5±4.0	60.0 58.5±3.7	0.0 -0.3±0.6	0.1969 0.196±0.003	0.750 0.745±0.008	0.250 0.245±0.010	1.305 1.309±0.011	14.54 14.53±0.09
θ_{60}C$_{6250}$ 58.2±7.7	1.60 1.62±0.03	12.0 12.5±0.5	60.0 58.6±6.6	60.0 60.2±6.6	0.0 -0.7±1.5	0.1969 0.192±0.008	0.750 0.728±0.027	0.250 0.247±0.023	1.305 1.303±0.028	14.54 14.52±0.18
$(M/R)_{hi}$ 60.3±7.7	1.68 1.61±0.11	11.0 10.4±1.0	60.0 34.9±13.1	20.8 42.7±16.2	0.0 2.1±4.8	0.2255 0.228±0.012	0.308 0.328±0.029	0.467 0.542±0.116	0.350 0.343±0.022	5.74 5.80±0.22
$(M/R)_{lo}$ 62.5±10.1	1.35 1.28±0.10	13.0 13.0±1.2	60.0 34.3±17.2	19.8 42.8±17.1	0.0 1.9±3.6	0.1534 0.146±0.018	0.293 0.299±0.027	0.471 0.522±0.086	0.423 0.411±0.022	5.75 5.78±0.15
$\beta_{hi}(M/R)_{hi}$ 53.8±10.7	1.68 1.70±0.02	11.0 11.2±0.3	60.0 61.9±6.5	64.0 61.6±6.8	0.0 -0.8±0.8	0.2255 0.223±0.004	0.778 0.762±0.016	0.219 0.211±0.016	1.235 1.241±0.016	14.53 14.47±0.09
ν_{400} 58.1±9.9	1.60 1.55±0.21	12.0 13.0±2.5	60.0 35.7±17.6	20.0 40.7±18.8	0.0 1.0±5.2	0.1969 0.181±0.034	0.296 0.292±0.051	0.470 0.525±0.155	0.373 0.361±0.035	3.83 3.88±0.22

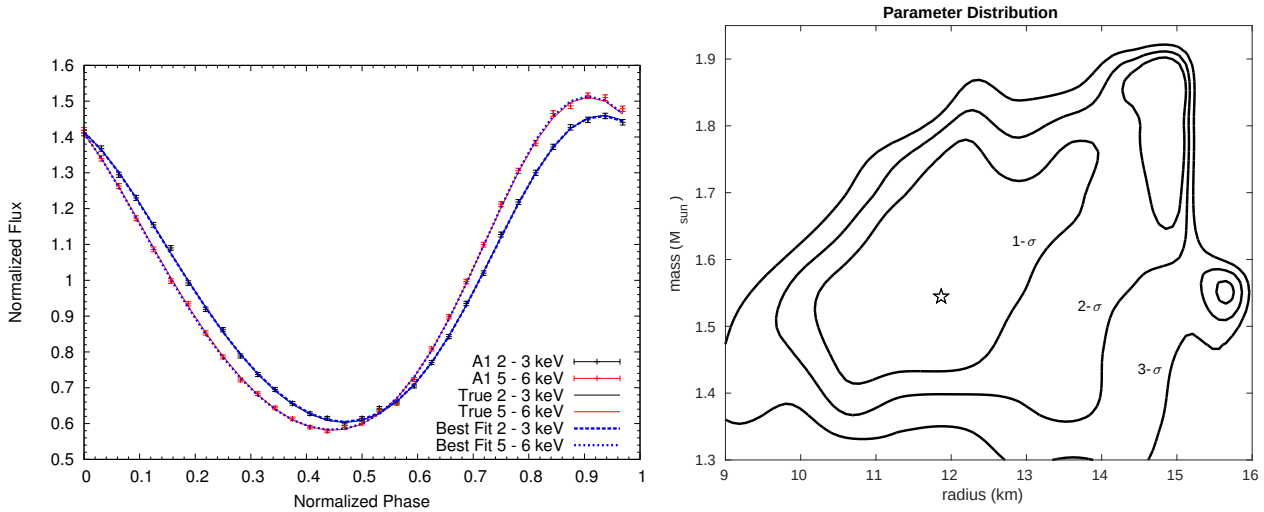

Fig. 1.— Left: Pulse profiles for model A in two energy bands. The solid black and red curves (labeled "True") correspond to the theoretical model A given in Table 1. The black and red points (labeled "A1") with error bars correspond to the addition of noise to model A with one Poisson realization assuming an average count-rate of 25,000 per phase bin. The blue dashed and dotted curves correspond to the best fit to model A1 with parameters given by the values in row "A1" in Table 3. This best fit has $\chi^2 = 58.2$ for 59 dof. Right: χ^2 contours in the $M-R$ plane for fitting synthetic data from model A1. The best-fit model is shown with a star and corresponds to the row labelled "1" in Table 3. The "true" values of mass and radius are 1.6 $M_\odot$ and 12 km. The contours show the 1-, 2-, and 3-$\sigma_{\rm degen}$ confidence regions. From this figure, it can be seen that the 1-$\sigma_{\rm degen}$ error region for radius corresponds to approximately 1.9 km, while the 1-$\sigma_{\rm degen}$ limit for mass spans 0.18 $M_\odot$. The contours for 1-, 2-, and 3-$\sigma_{\rm degen}$ correspond to values of $\Delta\chi^2$ (above $\chi^2_{\rm min}$) of 2.3, 6.2, and 11.8 respectively.

The best fit model (star) and contours of constant $\chi^2 = \chi^2_{\rm min} + \Delta\chi^2$ (black curves) for the A1 dataset are shown in the $M-R$ plane in Figure 1. The contours correspond to values of $\Delta\chi^2 = 2.3, 6.2$, and 11.8, corresponding to 1-, 2-, and 3-$\sigma_{\rm degen}$ confidence levels respectively in mass-radius space for two free parameters, M and R. This plot uses the profile likelihood (Murphy & Van Der Vaart 2000) to eliminate the nuisance parameters i, θ, and ϕ; a grid is defined in M and R, and for each (M, R) grid point, Ferret finds the lowest χ^2 within each grid cell, allowing for any values of i, θ, and ϕ. The 1-$\sigma_{\rm degen}$ confidence region (calculated by finding the maximum and minimum values on the 1-$\sigma_{\rm degen}$ contour and dividing by 2) gives an approximate 1-$\sigma_{\rm degen}$ error of 1.9 km for the radius and a 0.18 $M_\odot$ error for the mass. Note that since the contours are not ellipses, these are only approximate 1-$\sigma_{\rm degen}$ limits.

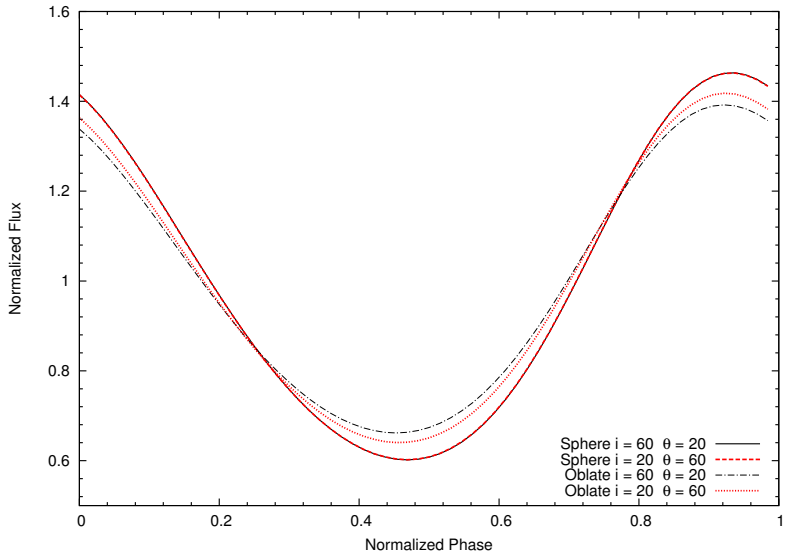

FIG. 2.— Normalized low-energy band pulse profiles for four neutron star models showing the observer inclination–spot co-latitude degeneracy for spherical and oblate stars. The solid black curve represents Model A, while the overlapping dashed red curve represents a star with the same parameters as Model A, but with inclination i and co-latitude θ swapped. The black dot-dashed curve represents Model AO, the oblate version of Model A. The red dotted curve represents a star with the same parameters as Model AO, again with inclination i and co-latitude θ swapped.

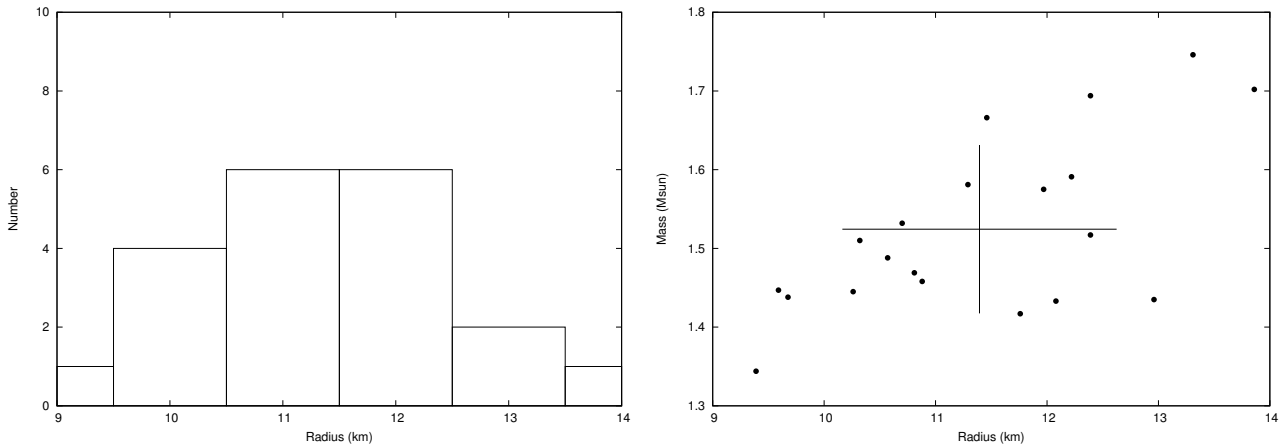

FIG. 3.— LEFT: Histogram of best-fit values of R for model A. RIGHT: Best fit values of M and R for model A. Each dot is one fit to a Poisson realization, as shown in Table 3. The black cross shows the average with 1-σ error bars.

Ferret was used to fit each of the 20 Poisson realizations of model A. The independent best fit results for the 20 different realizations are shown in Table 3. The average and standard deviation for each parameter are displayed at the bottom of Table 3, and also appear in the second line of Table 2. For the individual angles i and θ, it can be seen from Table 3 that the determinations of these angles are very poor. This is due to the already well-known degeneracy (for spherical stars; Poutanen & Gierliński 2003) which occurs since the equations for light-bending and the Doppler effect only depend on the combinations $\sin i \sin\theta$ and $\cos i \cos\theta$ and not on their individual values. The $i - \theta$ degeneracy can be seen in Figure 2, where the normalized low-energy band pulse profiles for model A, and another model with i and θ swapped (both labelled "Sphere") are indistinguishable. For many of the Poisson realizations, the best-fit values for i and θ shown in Table 3 are swapped from their true values, and as a result, the average and standard deviations for the individual angles are really not meaningful, except to illustrate that it is the trigonometric combinations of the angles that can be reliably determined. However, since there are independent methods for constraining i and θ through optical (Wang et al. 2013) or gamma-ray (Venter et al. 2012) observations, it is still useful to discuss these two angles separately.

To ensure that a suitable number of independent realizations of the data were made, a histogram of the best-fit radii is plotted in Figure 3. The resulting histogram is approximately Gaussian (70% are within 1 standard deviation, 95% within 2 standard deviations, 100% within 3 standard deviations), indicating that 20 trials are sufficient to illustrate general trends. A scatter plot, also shown in Figure 3, shows the values of mass and radius (red crosses) for the 20 different random realizations of the data. The large black cross indicates the average and standard deviation values for the mass and radius fits.

As previously stated, parameter degeneracy plays a major role in all aspects of uncertainty in our results, even the

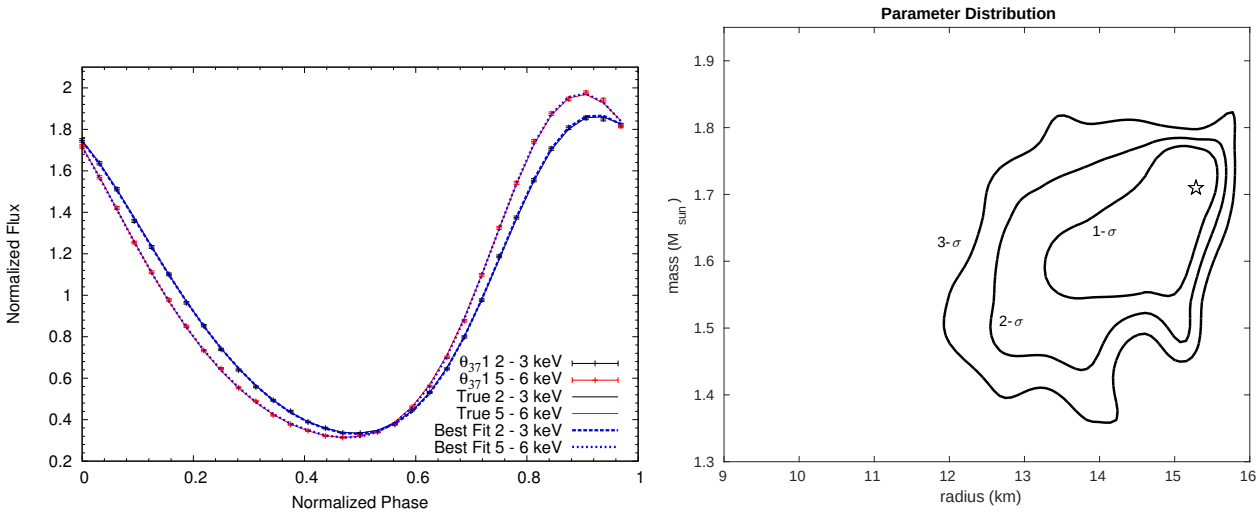

FIG. 4.— LEFT: Pulse profiles for model θ_{37} in two energy bands. RIGHT: χ^2 contours in the $M-R$ plane for fitting synthetic data from model $\theta_{37}1$. Symbols and lines have the same meaning as in Figure 1. The best fit values for mass and radius for this Poisson realization are 1.71 $M_\odot$ and 15.29 km, with $\chi^2 = 58.7$ for 59 dof. The approximate 1-$\sigma_{\rm degen}$ limits for mass and radius are 0.11 $M_\odot$ and 1.1 km.

standard deviation. By testing 20 different realizations of the data for each model (a relatively small number), we still find that some models converge on the true parameter values while others do not. The models with stronger parameter degeneracy have larger standard deviations on the parameters. This provides an initial assessment of which types of models and pulse profile shapes have stronger degeneracies. Even with very small error bars on the pulse profile, there would still be relatively large standard deviations over many models due to the inherent degeneracies.

3.2. *Effect of Changing Spot Co-latitude*

In this section, we examine the effect of changing the hotspot's co-latitude θ on the accuracy and precision of the pulse profile fits. By changing θ while keeping M, R, and i constant, we alter the projected velocity β and the approximate pulse amplitude Amp (see Table 1). Models A (with $\theta = 20°$), θ_{37}, and θ_{60} have increasing values of θ, β, and Amp. Due to the i-θ degeneracy, this is also equivalent to keeping θ fixed and varying i.

The $\theta_{37}1$ pulse profile (i.e., for the first Poisson realization of the θ_{37} model) is shown in the left panel of Figure 4 and the constant χ^2 contours are shown in the right panel. This case is one of the "outlier" pulse shapes (of the 20 Poisson realizations) with a very large best-fit R that only includes the true R at the 3-$\sigma_{\rm degen}$ level. The 1-$\sigma_{\rm degen}$ error regions for the $\theta_{37}1$ model (from the right panel of Figure 4) are smaller than the A1 model, close to 1.1 km for R and 0.11 $M_\odot$ for M. The 1-$\sigma_{\rm degen}$ regions for this particular Poisson realization are somewhat larger than the standard deviation 1-σ computed for the ensemble of θ_{37} models.

Similarly, the $\theta_{60}1$ pulse profile (first Poisson realization of the θ_{60} model) is shown in the left panel of Figure 5 and the constant χ^2 contours are shown in the right panel. The 1-$\sigma_{\rm degen}$ error regions for the $\theta_{60}1$ model are 0.7 km for R and 0.06 $M_\odot$ for M, larger than the ensemble standard deviation 1-σ by a factor of about 6.

We find a strong trend as θ increases: the average M and R fit values are closer to the true values and the standard deviation is smaller. For example, σ in M decreases from 0.11 to 0.08 to 0.01 $M_\odot$ for models Models A, θ_{37}, and θ_{60}. This improvement in fitting accuracy is due to models θ_{37} and θ_{60} having a larger Amp and β than model A. This is consistent with the general trend seen by Lo et al. (2013) in individual model fits.

During an X-ray burst, θ is expected to change. This set of models roughly approximates this process, which has been explored in detail by Mahmoodifar & Strohmayer (2016). If a series of *different* burst oscillation pulse shapes are found for a neutron star during the *same* burst, future work could include a simultaneous multi-epoch fit, allowing θ (and spot size ρ) to be dependent on epoch.

3.3. *Effect of Oblateness*

Rotation alters the shape of a neutron star, making it an oblate spheroid. Although the change in shape is small for stars spinning at rates seen in accreting systems, the alteration in the star's shape changes the positions on the neutron star's surface for which photons can reach the observer (Morsink et al. 2007), leading to large changes in the pulse profile. As an example, a 1.6 $M_\odot$ neutron star with an interior given by the APR equation of state (Akmal et al. 1998) has an equatorial radius of 11.7 km and a ratio of polar to equatorial radii of 0.93. This oblate geometry has a larger

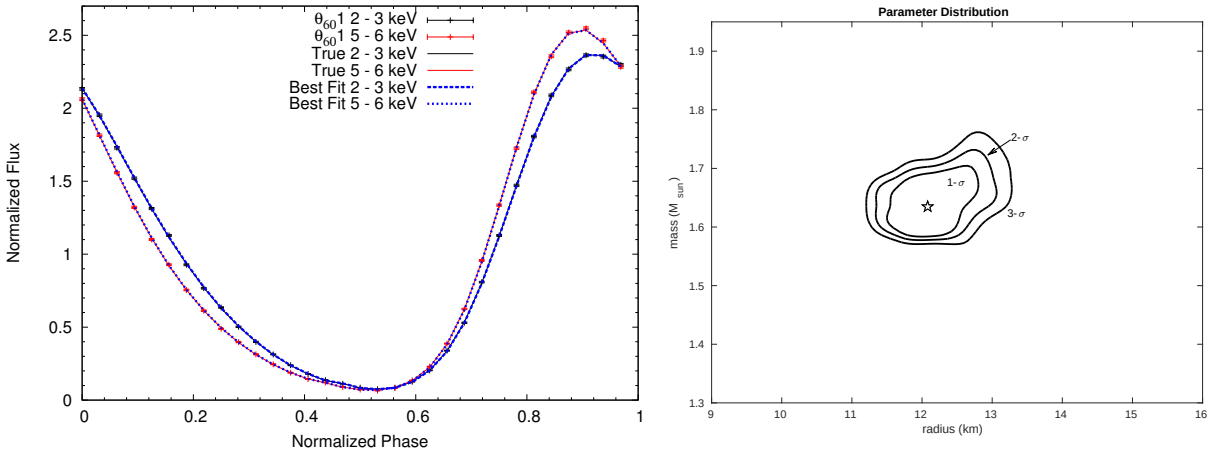

FIG. 5.— LEFT: Pulse profiles for model θ_{60} in two energy bands. RIGHT: χ^2 contours in the $M-R$ plane for fitting synthetic data from model $\theta_{60}1$. The symbols have the same meaning as in Figure 1. The best fit values for mass and radius for this Poisson realization are $M = 1.63\,\mathrm{M}_\odot$ and $R = 12.1$km, with $\chi^2 = 49.6$ for 59 dof. The approximate 1-$\sigma_{\rm degen}$ limits for mass and radius are 0.06 $\mathrm{M}_\odot$ and 0.7 km.

effect on the pulse profile than other effects due to rotation, such as frame-dragging, as has been discussed in detail in, e.g., Morsink et al. (2007).

In this section, we investigate the effect of the oblate shape on the accuracy and precision of pulse profile fitting and reproducing the input parameters. To do this we construct oblate versions of two models, A and θ_{37}. The oblate versions, designated with the letter "O" are constructed so that the radius of the star at the location of the spot is the same as for the corresponding spherical model. This means that the values of the star's compactness M/R and the projected velocity β are the same for the spherical and oblate models. As a result of this definition for the radius, the star's equatorial radius is larger than that listed in Table 1. The general trend is for the oblate neutron star's pulse profile to have a smaller pulse amplitude than the spherical neutron star with the same parameters (when i and θ are in the same hemisphere), since the visibility condition makes it easier to see the far side of the neutron star when the star is oblate (Cadeau et al. 2007; Morsink et al. 2007). In cases where the spot is visible for all phases (as in both models A and θ_{37}), consider the emission when the spot and the observer have the same azimuthal angle. At this moment, the light from the spherical star is emitted close to the normal to the surface. For the oblate star, the light is emitted in the same direction in space, but it is at an angle further from the surface's normal, due to the tilt of the surface. Since the intensity of light is proportional to the cosine of the angle between the original direction of emission and the normal to the surface, the spherical star's spot appears brighter at this phase. The opposite is true when the spot is on the opposite side of the star from the observer. In this case the light is emitted close to the tangent to the surface for the spherical star, and closer to the normal for the oblate star, making the spot appear dimmer for the spherical star at this phase. The overall effect is to create a less modulated pulse profile for the oblate star.

In model AO, the neutron star's equatorial radius is 12.7 km, while the radius at the spot's latitude is 12 km. As a comparison with the spherical model, the low-energy band pulse profiles for both AO and A are plotted as black dot-dashed and solid curves respectively in Figure 2. The high energy band (not shown for clarity) has a similar decrease in modulation. In model θ_{37}O the equatorial radius is 12.5 km, and has a lower pulse amplitude than model θ_{37}. Since the pulse profiles and $M-R$ confidence regions for models AO and θ_{37}O look quite similar to those for models A and θ_{37}, they are not shown here. The mean and standard deviation of the best fit parameters for the 20 realizations for each oblate model are given in Table 2. For most parameters, the σ and accuracies are smaller for the oblate case than for the fiducial spherical case (we also found that the 1-$\sigma_{\rm degen}$ regions were smaller by about a half). A similar improvement in the precision of the results was also noted by Miller & Lamb (2015), however since they only considered one Poisson realization of the data they could not rule out the possibility that this was due to a statistical fluctuation. In our case, since we are comparing a sample of Poisson realizations for each model, the increase in precision and accuracy is most likely due to the properties of the oblate model pulse profiles.

The improvement in the accuracy and precision in the determination of most of the parameters for oblate models is most likely due to the partial lifting of the degeneracy between i and θ. For an oblate star, the direction that the normal to the surface points in depends on the shape of the star which is a function of θ, but is independent of i. This introduces small differences between the normalized pulse profiles for models with i and θ switched, while for spherical stars, the normalized pulse profiles are the same when the angles are switched. This is illustrated in Figure 2 where the two oblate models representing the swapped inclination and spot angles, shown with black dot-dashed and red dotted curves, are clearly different from each other. However, since they are still fairly similar to each other, there is still a partial degeneracy when the angles are swapped.

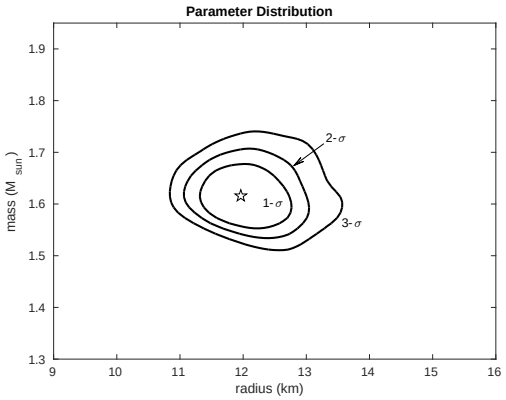

FIG. 6.— χ^2 contours in the $M-R$ plane for fitting synthetic data from model $\theta_{60}C_{6250}1$. The symbols have the same meaning as in Figure 1. The best fit values for mass and radius for this Poisson realization are $M = 1.62\,\mathrm{M}_\odot$ and $R = 12.0\,\mathrm{km}$, with $\chi^2 = 53.8$ for 59 dof. The approximate 1-$\sigma_{\rm degen}$ limits for mass and radius are 0.06 $\mathrm{M}_\odot$ and 0.8 km.

There is a similar but much smaller lifting of degeneracy for large hot spots on spherical stars. A large spot extends over a range of θ values while the observer's inclination i is just one fixed value, so swapping the two angles if the spot is large will not yield the same pulse shape. However, the magnitude of the effect is much smaller than the magnitude of the change that occurs when the angles are swapped on an oblate star.

We also tested the effect of using the wrong shape model by using the oblate data corresponding the 20 AO models as input and fitting them with pulse shapes for spherical stars. In this case, χ^2 increased a small amount but the best-fit values for M and R became very inaccurate. For the set of fits to the Poisson realizations, the average fit returned $\chi^2 = 64.7$, $M = 1.91\,\mathrm{M}_\odot$, and $R = 12.3\,\mathrm{km}$ (which should be compared with the values in the 4th row of Table 2). The average M and R fits are accurate to 19.4% and 2.5%, respectively. Similar results were found by Miller & Lamb (2015). In practice, this is not a problem, since we know that the stars are actually oblate, so the correct shape can be used. The spherical shape model is computationally somewhat cheaper, hence it is used for the other models tested in this paper. Furthermore, the oblate shape model used is still an approximation. It would be worthwhile, in future research, to compare the slightly different shape models that have been used in this paper with other models used by other groups (such as Bauböck et al. 2013; Miller & Lamb 2015).

3.4. *Effect of Photon Count-Rate*

The average photon count-rate per phase bin of 25000 is the expected best case scenario, assuming typical burst flux, large detector effective area, and long observation time (Feroci et al. 2012), and similar to the count-rate used in other pulse profile analyses (Lo et al. 2013; Psaltis et al. 2014). The pulse profile is noisier when fewer counts are detected. We tested the effect of statistical noise on model θ_{60} by generating different Poisson realizations, assuming a reduced count-rate, of the same theoretical pulse profile. Comparing model θ_{60} with the reduced count-rate model, $\theta_{60}C_{6250}$, allows us to explicitly compare the degeneracy-related uncertainty versus error. The average counts per bin are 25000 and 6250 for models θ_{60} and $\theta_{60}C_{6250}$, respectively.

The pulse profile for model $\theta_{60}C_{6250}$ looks very much like that for model θ_{60} but with larger errors, so it is not shown. The χ^2 contours for model $\theta_{60}C_{6250}1$ are shown in Figure 6. The size of the 1-$\sigma_{\rm degen}$ contours in Figures 5 and 6 are very similar, suggesting that the confidence contours for one particular mock observation are mainly dominated by the parameter degeneracy.

From Table 2, the standard deviations of the parameters for the low count-rate model are $\simeq 2-3$ times as large as those of model θ_{60}. The factor of two is expected purely on statistics (from the factor 4 reduction in counts). The actual degradation is somewhat worse, likely because partial degeneracy between the parameters makes the parameter determination degrade more rapidly than the errors. Model $\theta_{60}C_{6250}$ still gives well-determined parameters even at the lower count-rate value. The accuracies in M and R determination are 2% and 4% respectively, for the low count-rate case.

3.5. *Effect of Compactness*

Changes in the compactness ratio (M/R) are expected to impact the quality of the fits, since the compactness affects the pulse amplitude Amp (evident in equation (1)). Increasing the compactness decreases Amp, which might be

expected to decrease the precision and/or accuracy of the fits. To test this effect, we generated models with differing values for the compactness ratio, but the same values of the projected spot velocity β.

We created two models similar to model A with a larger and smaller compactness ratio, $(M/R)_{hi}$ and $(M/R)_{lo}$, and generated a set of pulse shapes with Poisson noise. The results of the fits can be seen in Table 2. All three models with the same value of β have similar accuracy and precision for most of the parameters. For instance, the accuracy of determining the mass ranges from 4 – 5%, while the precision ranges from 6 – 7%, so the change in compactness does not greatly affect the fits. However, while there is little change in the precision of the radius measurement (9 – 10%), in the case of the low compactness model the accuracy was improved.

As a further test on the effect of compactness, model $\beta_{hi}(M/R)_{hi}$ was created to compare with the high spot co-latitude model θ_{60}. The parameters were the same except that θ was adjusted to give the same β as model θ_{60}. This also necessitated that Amp was somewhat lower than model θ_{60} (see equations (1) and (2)). Previously for model θ_{60}, the accuracy and precision for both M and R were better than 1%; by fitting model $\beta_{hi}(M/R)_{hi}$, we found that increasing the compactness degrades both the accuracy and precision in M and R, but only so that they range from 1-3%.

From these results, given a value of β, it appears that the accuracy and precision to which the other parameters can be determined does not strongly depend on the compactness ratio of the star. Thus, our fit results would not be significantly different if we had chosen a different M and R for the fiducial model A.

3.6. *Effect of Spin Frequency*

The last parameter we adjusted to be different from the fiducial model was $\nu_{\rm spin}$. We compared model A ($\nu_{\rm spin}$ = 600 Hz) with a 400 Hz model, labelled ν_{400}. Model ν_{400} has all parameters the same as model A except $\nu_{\rm spin}$, and hence a different projected spot velocity (β) (see Table 1). This model has parameter values very similar to the "low inclination" model presented by Lo et al. (2013).

The fits results are summarized in Table 2. The precision (standard deviations) and accuracy (difference of mean and true values) are worse for model ν_{400}, as expected. For M, R, M/R, and Amp, the precisions are about twice as large for the 400 Hz model as for the standard 600 Hz model. However, β has the same absolute precision for the 400 Hz model as for the 600 Hz model.

3.7. *Effect of Incorrect Model or Additional Free Parameters*

Along with the sets of trials summarized in the previous subsections, we tested some of the model assumptions used in the fits. These tests include the assumptions we have made about the spectral emissivity, spot size, presence of a constant background flux, and the star's shape. Errors are not reported for the tests of this subsection since fits were only carried out on the A1 synthetic data set, the first Poisson realization of the Model A theoretical curve.

We first tested the effect of using the wrong atmosphere model in the fits. To do this, we used the synthetic data corresponding to the A1 model (see Figure 1) as input to Ferret. The synthetic data were generated using the Hopf limb-darkening function, but Ferret was run assuming a perfectly isotropic blackbody spectrum. The algorithm was unable to find a good fit after 300 generations (about double the number of generations normally required) and the best fit had $\chi^2 = 316.6$ for 59 dof. The best-fit M and R were 1.21 $M_\odot$ and 11.3 km, which are quite inaccurate (24.4% for M, 5.8% for R). As was seen by Lo et al. (2013), one does not get a good fit using an incorrect atmosphere model.

We next tested fitting a variable spot size model to a pulse profile that was generated with an infinitesimally small spot. The infinitesimal size of the spot is a necessary assumption due to the large number of trials carried out in this work. In order to model a larger spot, the spot has to be cut into a number of segments, and separate computations of the deflection angles must be computed, increasing the run time linearly with the number of segments. The pulse shape for a large spot tends to have a lower pulse amplitude than a small spot centred at the same latitude, due to the effect of averaging over many latitudes. As a test of this assumption, we used the A1 data set as input to the Ferret algorithm, but added the angular radius ρ of the spot as a free parameter that was allowed to vary between 1°and 60°. The resulting best fit has $\chi^2 = 59$ (for 58 dof) and converged on an angular radius of $\rho = 1°$, the smallest angle allowed. The best-fit M and R for this case are 1.51 $M_\odot$ and 11.6 km. While these values are less accurate (5.6% and 3.3%, respectively) than the best-fit value shown in the first row of Table 3 corresponding to a fit with an infinitesimal spot, they are within the 1-σ limits computed for the set of model A fits. Since using multiple spot segments leads to a significantly longer computation time, it was not feasible to use variable spot sizes in this paper. However, in future work when real data are being fitted, it will be necessary to allow the spot size to vary, and to compute the observed flux from multiple spot segments.

We also tested the effect of allowing the temperature of the spot to be a free parameter in the fitting program. The theoretical model was constructed with a spot temperature of 2 keV (as measured in the comoving frame at the star's surface). The Ferret program was used to fit the data with the addition of a local temperature parameter that was allowed to vary between 1 and 3 keV. The best fit returned $\chi^2 = 57.9$ (for 58 dof), $M = 1.67\,M_\odot$, $R = 11.8$ km, and a temperature of 2.1 keV. The accuracies in M and R are 4.8% and 1.7%, respectively. Although these M and R values are within the 1-σ limits for this data set, the addition of the extra parameter does degrade the accuracy. This is most likely due to the degeneracy introduced since we measure the redshifted temperature with the light curves. For the case of real data, it would be important to allow the temperature to vary in the fits, since the temperature of the spot in the comoving frame is unknown. Furthermore, making use of multiple energy bands would improve the accuracy of the temperature measured at infinity.

The effect of an unknown background flux has been investigated in detail by Lo et al. (2013), and for comparison we test it for one simple case. The synthetic data set A1 was constructed without a background count rate. We used this as input and allowed Ferret to add two new parameters corresponding to a constant (in time) background flux in each energy band. These background values were allowed to vary between 0 and 1 (recall that our pulse profiles are normalized to 1). The resulting fit had $\chi^2 = 57.7$ (for 57 dof), and the best-fit M and R were $1.52\,M_\odot$ and 11.7 km. M and R were accurate to 5.0% and 2.5%, respectively. It is possible to add a more realistic background model by adding emission at a lower temperature from the rest of the star (Lo et al. 2013; Elshamouty et al. 2016) or by adding light scattered from the disk (Morsink & Leahy 2011), but at this point it is not obvious what the most realistic model would be. Real data are also likely to contain non-negligible emission from the surface of the neutron star outside the spot region and from the accretion disk.

The effect of these tests on the best-fit values of M and R are summarized in Table 4. These tests highlight the importance of having as realistic an atmosphere model as possible, since fitting with the wrong model drastically affected the quality and accuracy of the fits. Thus, this method provides a sensitive test for atmosphere models. Adding a free parameter for the temperature, unknown background flux, or spot size did not significantly detract from the accuracy in M and R, and our code was able to replicate the additional parameter quite well. Ideally, for real data we would allow these parameters to also vary in the fits, even if it is not feasible to do so in the present study.

4. DISCUSSION AND CONCLUSION

We calculated pulse profiles for simulated thermonuclear burst oscillations from a rapidly rotating neutron star as would be detected by a next-generation X-ray timing observatory, such as ASTROSAT, NICER, or LOFT. The input neutron star parameters include mass M, radius R, emitting spot co-latitude θ, and observer inclination i. We created 20 Poisson realizations of each test pulse profile, which allows us to determine the errors in the derived model parameters, focusing on M and R. The resulting pulse profiles were fitted with Ferret to analyze how well the input parameters could be recovered (in both standard deviation and degeneracy-related uncertainty).

We find that the best-determined parameter is the projected velocity β of the spot along the observer's line-of-sight. The next best-determined are M and R. Compactness (M/R), $\sin i \sin\theta$, and the pulse amplitude Amp are also well-determined, but $\cos i \cos\theta$, i, and θ are poorly determined. It is clear that more rapidly rotating neutron stars produce pulse profiles with strong harmonics, which yield better constraints on their parameters. Also, neutron stars viewed at larger i with spots at a larger θ produce more modulated pulse profiles and therefore better constraints. However, we also found that count-rate is important, so the best targets will be those systems which have the optimal combination of large $\nu_{\rm spin}$, large θ and i, and bright pulsations. For our best cases presented here, M and R were determined to 1% accuracy and precision, but even for many of the less optimum cases M and R were determined to ~5% accuracy and precision, which is a very valuable result.

We carried out a number of parameter comparison tests, to see how different input parameters can affect the constraints on M and R. The more asymmetry and larger pulse amplitude, the better the accuracy and precision in M and R from pulse profile fitting. The asymmetry in the pulse profile is controlled by β; a larger β results from increases in θ, i, and $\nu_{\rm spin}$. A larger θ or i also gives a larger Amp. Compactness (M/R) has a small effect on parameter determination; for more compact stars it is somewhat more difficult to determine parameters. This is caused by the increased visibility of the surface by the observer and the resulting decreased pulse amplitude. Including oblateness in the pulse profile model improves the accuracy and precision of M and R determinations. This is due to a reduced $i-\theta$ degeneracy in the oblate models.

Photon count-rate has a critical effect on parameter determination. On simple grounds, increasing the counts by a factor of four reduces the error by a factor of two. However, in practise we found a factor of three improvement in parameter determination. The extra gain is likely the result of reduced parameter degeneracy for data with smaller errors.

We assumed the spot temperature to be a known definite quantity. In practice, this is determined by a spectral fit to the data using multiple energy bands. We calculated the simulated pulse profiles using only two energy bands, so we have underestimated the uncertainties in M and R that enter through the uncertainty of the temperature. In principle, it is not difficult to use more energy bands (e.g., Lo et al. 2013). However, there is a trade off that the signal-to-noise in each energy band falls as more, and thus smaller, energy bands are used. Determining an optimal number of energy bands that allows a determination of the spectrum while providing enough statistics to constrain the star's properties should be the topic of a future study.

We also took the background count rate to be negligible in comparison with the count rate from the neutron star hot spot. For a real observation, the background comes from three sources: instrument background, sky background, and source contribution to background. Instrument background depends strongly on the instrument design and mode of observation. Source background can consist of emission from the surface of the neutron star outside of the spot region and from the accretion disk. However, subtracting the persistent pre-burst emission from the X-ray burst flux is not an appropriate source-background subtraction. As shown in Worpel et al. (2013), the persistent emission can be different during the burst, possibly due to an increased accretion rate onto the neutron star. A method such as weighted-photons Bayesian Blocks (Worpel & Schwope 2015) should be incorporated into the analysis pipeline to subtract the appropriate level of source-background emission.

A limitation of our study is that we did not include the effect of frequency drift, which is normally seen during the rise of an X-ray burst (Watts 2012). When dealing with real data, it would be important to either remove sections of the data where frequency drift is observed, or to use a reliable method for modelling the pulse shape with phase offsets

to account for the drift. An additional complication for the sources that also have accretion-powered pulsations is that the pulsations may contaminate the X-ray burst oscillations. At this time these are open problems associated with modelling X-ray burst oscillations.

In this paper we have demonstrated that an evolutionary optimization and parameter search methods can be used to effectively constrain the mass and radius of a neutron star with pulsed emission. We focussed on burst oscillations seen in Type I X-ray bursts, however the routines can also be used to model accretion-powered pulsations or rotation-powered pulsations once the appropriate spectrum and beaming functions have been changed. In particular, the rotation-powered pulsars that will be studied by NICER (see Ozel et al. 2015) can be analyzed using the methods discussed in this paper.

This research was supported by grants from the Natural Sciences and Engineering Research Council of Canada to J.D.F., D.A.L., and S.M.M. A.L.S. acknowledges a travel grant from the LKBF (Leids Kerkhoven-Bosscha Fonds). We thank Arash Bahramian, Frederick Lamb, M. Coleman Miller, Adam Rogers, and Christoph Weniger for useful discussions. We also thank the LOFT Dense Matter science working group for early code comparisons to validate the modelling.

REFERENCES

Akmal, A., Pandharipande, V. R., & Ravenhall, D. G. 1998, Phys. Rev. C, 58, 1804
AlGendy, M., & Morsink, S. M. 2014, ApJ, 791, 78
Artigue, R., Barret, D., Lamb, F. K., Lo, K. H., & Miller, M. C. 2013, MNRAS, 433, L64
Arzoumanian, Z., Gendreau, K. C., Baker, C. L., et al. 2014, Proc. SPIE, 9144, 914420
Bauböck, M., Berti, E., Psaltis, D., & Özel, F. 2013, ApJ, 777, 68
Bauböck, M., Psaltis, D., Özel, F. 2015, ApJ, 811, 144
Beloborodov, A. M. 2002, ApJ, 566, L85
Bogdanov, S. 2013, ApJ, 762, 96
Cadeau, C., Morsink, S. M., Leahy, D. A., & Campbell, S. S. 2007, ApJ, 654, 458
Chandrasekhar, S. 1960, Radiative Transfer (Dover Publications, Inc.)
Elshamouty, K. G., Heinke, C. O., Morsink, S. M., Bogdanov, S. & Stevens, A. L. 2016, ApJ, 826, 162
Feroci, M., Stella, L., van der Klis, M., et al. 2012, Experimental Astronomy, 34, 415
Fiege, J. D. 2010, Qubist User's Guide: Optimization, Data-Modeling, and Visualization with the Qubist Global Optimization Toolbox for MATLAB, 4th edn., nQube Technical Computing Corporation
Fiege, J. D., McCurdy, B., Potrebko, P., Champion, H., & Cull, A. 2011, Med. Phys. 38, 5217
Franzmann, E. 2014, MSc thesis, Univ. of Manitoba, URL: http://mspace.lib.umanitoba.ca/handle/1993/23830
Galloway, D. K., Muno, M. P., Hartman, J. M., Psaltis, D., & Chakrabarty, D. 2008, ApJS, 179, 360
Ghosh, P., & Lamb, F. K. 1979, ApJ, 232, 259
Goldberg, D. E. 1989, Genetic Algorithms in Search, Optimization, and Machine Learning (Addison-Wesley Publishing Company, Inc.)
—. 2002, The Design of Innovation: Lessons from and for Competent Genetic Algorithms (Kluwer Academic Publishers)
Heinke, C. O., Rybicki, G. B., Narayan, R., & Grindlay, J. E. 2006, ApJ, 644, 1090
Holland, J. H., 1975, Adaptation in Natural and Artificial Systems (Ann Arbor, MI: University of Michigan Press)
Leahy, D. A. 2004, ApJ, 613, 517
Leahy, D. A., Morsink, S. M., Chung, Y.-Y., & Chou, Y. 2009, ApJ, 691, 1235
Leahy, D. A., Morsink, S. M., & Chou, Y. 2011, ApJ, 742, 17
Lo, K. H., Miller, M. C., Bhattacharyya, S., & Lamb, F. K. 2013, ApJ, 776, 19
Mahmoodifar, S., & Strohmayer, T. 2016, ApJ, 818, 93
Miller, M. C., & Lamb, F. K. 1998, ApJ, 499, L37
Miller, M. C., & Lamb, F. K. 2015, ApJ, 808, 31
Morsink, S. M., & Leahy, D. A. 2011, ApJ, 726, 56
Morsink, S. M., Leahy, D. A., Cadeau, C., & Braga, J. 2007, ApJ, 663, 1244
Murphy, S. A., & Van Der Vaart, A. W. 2000, Journal of the American Statistical Association, 95, 450
Ozel, F., Psaltis, D., Arzoumanian, Z., Morsink, S., & Baubock, M. 2015, arXiv:1512.03067
Pechenick, K. R., Ftaclas, C., & Cohen, J. M. 1983, ApJ, 274, 846
Poutanen, J., & Gierliński, M. 2003, MNRAS, 343, 1301
Poutanen, J., Ibragimov, A., & Annala, M. 2009, ApJ, 706, L129
Psaltis, D., & Özel, F. 2014, ApJ, 792, 87
Psaltis, D., Özel, F., & Chakrabarty, D. 2014, ApJ, 787, 136
Rogers, A., & Fiege, J. D. 2011, ApJ, 727, 80
Rogers, A., & Fiege, J. D. 2012, ApJ, 759, 27
Rogers, A., Safi-Harb, S., & Fiege, J. 2015, IAU General Assembly, 22, 2257588
Singh, K. P., Tandon, S. N., Agrawal, P. C., et al. 2014, Proc. SPIE, 9144, 91441S
Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2010, ApJ, 722, 33
Stevens, A. L. 2013, MSc thesis, Univ. of Alberta, ISBN: 9780494960219
Strohmayer, T. E., Zhang, W., Swank, J. H., White, N. E., & Lapidus, I. 1998, ApJ, 498, L135
Venter, C., Johnson, T. J., & Harding, A. K. 2012, ApJ, 744, 34
Wang, Z., Breton, R. P., Heinke, C. O., Deloye, C. J., & Zhong, J. 2013, ApJ, 765, 151
Watts, A. L. 2012, ARA&A, 50, 609
Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, Rev. Mod. Phys. 88, 021001
Wijnands, R., & van der Klis, M. 1998, Nature, 394, 344
Worpel, H., Galloway, D. K., & Price, D. J. 2013, ApJ, 772, 94
Worpel, H., & Schwope, A. D. 2015, A&A, 578, A80

TABLE 3
Summary of Best Fits for Model A

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	20.0	0.0000	0.1969	0.296	0.470	0.373	5.741
1	58.2	1.575	11.97	48.7	25.0	0.0071	0.1943	0.318	0.598	0.347	6.120
2	60.8	1.517	12.39	19.9	57.3	-0.0011	0.1808	0.286	0.508	0.362	5.577
3	58.6	1.344	9.39	39.4	34.4	0.0021	0.2114	0.358	0.638	0.357	5.565
4	71.7	1.532	10.70	21.7	59.8	0.0000	0.2114	0.319	0.467	0.383	5.658
5	63.9	1.666	11.46	71.8	17.6	-0.0062	0.2147	0.287	0.297	0.426	5.471
6	56.7	1.445	10.26	36.7	36.0	0.0064	0.2080	0.351	0.649	0.350	5.932
7	64.6	1.469	10.81	37.2	34.4	0.0074	0.2007	0.342	0.657	0.345	6.007
8	53.9	1.694	12.39	68.7	17.3	-0.0047	0.2019	0.277	0.347	0.404	5.584
9	42.3	1.417	11.76	42.8	27.5	0.0044	0.1779	0.313	0.651	0.338	5.773
10	67.0	1.510	10.32	21.9	60.6	-0.0007	0.2161	0.324	0.456	0.388	5.583
11	57.1	1.581	11.29	21.7	58.1	0.0026	0.2068	0.314	0.491	0.372	5.828
12	68.9	1.433	12.08	29.1	39.9	0.0062	0.1752	0.311	0.671	0.331	5.869
13	58.4	1.458	10.88	35.6	35.8	0.0075	0.1979	0.341	0.659	0.345	5.994
14	63.7	1.702	13.86	15.9	67.6	-0.0057	0.1813	0.254	0.366	0.390	5.541
15	58.7	1.447	9.59	38.1	36.9	0.0067	0.2228	0.371	0.629	0.360	6.006
16	45.5	1.746	13.31	18.3	63.1	0.0015	0.1937	0.280	0.429	0.375	5.983
17	51.8	1.435	12.96	30.7	36.3	0.0086	0.1635	0.302	0.693	0.323	6.009
18	53.1	1.488	10.57	38.4	34.2	0.0076	0.2079	0.349	0.648	0.348	6.076
19	72.0	1.591	12.22	19.9	59.0	0.0003	0.1923	0.292	0.485	0.367	5.727
20	58.1	1.438	9.68	38.5	36.1	0.0060	0.2195	0.367	0.632	0.359	5.962
Ave	59.3	1.524	11.39	34.7	41.8	0.0028	0.1989	0.318	0.549	0.363	5.813
Std	7.6	0.107	1.23	15.0	15.2	0.0046	0.0159	0.032	0.120	0.025	0.203

The summary of fits for the other models are available online only.

TABLE 4
Summary of Tests on A1 Data

Model	χ^2/dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	20.0	0.0000	0.1969	0.296	0.470	0.373	5.741
A1 Best Fit	58.2/59	1.575	11.97	48.7	25.0	0.0071	0.1943	0.318	0.598	0.347	6.120
Wrong Atmosphere	316.6/59	1.210	11.29	40.2	39.2	0.0104	0.1583	0.408	0.592	0.495	7.004
Spot Size Varies	59.0/58	1.508	11.57	33.3	37.0	0.0090	0.1925	0.331	0.667	0.338	6.137
Temperature Varies	57.9/58	1.668	11.81	23.1	55.1	0.0057	0.2086	0.322	0.526	0.364	6.266
Background	57.7/57	1.522	11.75	45.8	26.7	0.0066	0.1913	0.321	0.623	0.344	6.045